WEST Search History

Hide Items | Restore | Clear | Cancel

DATE: Thursday, June 30, 2005

Hide?	Set Name	Query	Hit Count
		DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ	
☐	L24	delayed same (target adj2 bandwidth)	4
☐	L23	L22 NOT l20	1
☐	L22	((wait or pause or delayed) adj2 time) same (target adj2 bandwidth)	8
☐	L21	L20 NOT l18	0
☐	L20	((wait or pause) adj2 time) same (target adj2 bandwidth)	7
☐	L19	20000419	0
☐	L18	(wait adj2 time) same (target adj2 bandwidth)	7
☐	L17	(timing adj2 module) same (data near5 transfer) same (target adj2 bandwidth)	2
		DB=PGPB,USPT; PLUR=YES; OP=ADJ	
☐	L16	20000419	0
☐	L15	(timing adj2 module) same (data near5 transfer) same (target adj2 bandwidth)	2
☐	L14	20000419	0
☐	L13	(wait adj2 time) same(byte near5 (count or counting)) same (target adj2 bandwidth)	2
☐	L12	20000419	3308
☐	L11	wait near5 (time or state) near5 (control or controlling or calculate or calculating or determine or determining)	5469
		DB=PGPB,USPT,USOC,EPAB,JPAB,DWPI,TDBD; PLUR=YES; OP=ADJ	
☐	L10	L9 and (target adj2 bandwidth)	0
☐	L9	20000419	4055
☐	L8	20000419	0
☐	L7	20000419	9
☐	L6	L5 and (target adj2 bandwidth)	33
☐	L5	(control or controlling or adjust or adjusting) near8 ((data or transfer or transmission or transfering) adj2 rate)	15858
☐	L4	l3 and (target adj2 bandwidth)	24
☐	L3	wait near5 (time or state) near5 (control or controlling or calculate or calculating or determine or determining)	6511
☐	L2	wait near5 (time or state) near5 (control or controlling)	4697
☐	L1	wait near5 (time or state)	60074

Previous Doc Next Doc Go to Doc#
First Hit Fwd Refs

Generate Collection

L7: Entry 2 of 9 File: USPT Sep 14, 2004

DOCUMENT-IDENTIFIER: US 6791943 B1
TITLE: Event-based technique for determining instantaneous cell bandwidth in a digital communications network

Application Filing Date (1):
19980227

Brief Summary Text (13):
The RM cells inserted into the backward-going ATM stream contain source-control information that instructs the source, node to either decrease or increase its transmission rate, depending on the current state of the network. That is, if additional bandwidth is available on the network, then a backward-going RM cell will indicate to the source that it may increase its transmission rate so as to fully utilize the available bandwidth. Conversely, if the bandwidth of the network is close to or above its maximum capacity, the backward-going RM cells may instruct an ATM source node to decrease its transmission rate so as to avoid cell loss caused by excessive congestion.

Detailed Description Text (43):
At step 904, the cell bandwidth determinator 402 counts the number of target cells that are received during the adaptive timebase. The cell bandwidth determinator 402 then determines the target cell bandwidth during the adaptive timebase at step 906 given the time interval and the counted number of received target cells.

Previous Doc Next Doc Go to Doc#

L7: Entry 7 of 9 File: USPT Jul 18, 2000

DOCUMENT-IDENTIFIER: US 6091777 A
TITLE: Continuously adaptive digital video compression system and method for a web streamer

Application Filing Date (1):
19980526

Brief Summary Text (31):
The novel method of the present invention contemplates compressing input color video data for moving images for a web stream to be transmitted over a transmission channel having a variable bandwidth to client users having receivers that may receive data input at different acceptable data rates. This novel method includes the steps of compressing an input color video data stream for moving images on the basis of a target bandwidth of a data transmission channel to provide a compressed output video data stream and transmitting the compressed output color video data stream to at least one receiver over the transmission channel. The actual bandwidth of the transmission channel is then determined as the compressed output video data stream was transmitted to at least one receiver. The compression of the input color video data stream is then continuously changed on a frame by frame basis in response to such determination whereby the transmission of a subsequent output color video data stream is commensurate with the available bandwidth of the transmission channel.

Brief Summary Text (32):
This novel method contemplates compressing the input color video data stream by detecting the amount of motion that is sensed in the input color video data stream and changing the rate of compression of said input color video data stream to provide an output color video data stream having a data rate commensurate with the available bandwidth of the transmission channel. This data rate by also be adjusted by adjusting the frame rate of the output compressed data stream so that the output data stream becomes commensurate with the available bandwidth of the transmission channel and

Detailed Description Text (28):
The input frame is linearly filtered row-wise and column-wise in four ways: Low-pass rows-low pass columns, low-pass rows-high -pass columns, high-pass rows-low-pass columns, and high-pass rows-high-pass columns. As the input frame is filtered it is also decimated by a factor of 2. Thus, the final filtered frames will each be one-fourth the sized of the input frame. Each of these subbands can now be treated as an independent stream and is processed further through motion detection to detect motion between successive frames and providing a dynamic priority to dynamically adjust the data rate of the output color video data stream to meet a variable target rate output based on the presently available network resources such as the available bandwidth of the transmission channel, quality and resolution desired by the client users, and receiver resources of the clients.

Detailed Description Text (45):
In order to keep the output data stream at desired level a target bandwidth is selected based on the transmission channel capacity and the receiver capability of

the least capable receiver of a client user. The noted thresholds T.sub.1 and T.sub.2 are automatically updated or changed on a frame by frame basis by monitoring the actual bandwidth of the transmission channel and generating a control signal from the difference between the actual bandwidth and the target bandwidth to vary or reset T.sub.1 and T.sub.2 as will be described in detail hereinafter.

Detailed Description Text (52):
Motion video may be considered as a sequence of images that may consist of different scenes involving low to high motion. In order to maintain a reasonably constant output data rate, the method of the instant invention provides an adaptive arrangement to adjust such output data rate to be commensurate with the available transmission channel bandwidth and the receiver resources of the client users that is based on the following two observations. The first is that motion and resolution in an image can be traded off since the human eye is not sensitive to details in a high motion picture sequence. The second is that higher resolution in an image is influenced by the higher subbands. That is for more details in an image, higher subbands need to be included in sending the video data stream.

Detailed Description Text (92):
The appropriate row in the Delta Bw Index Table is chosen in the following manner. One compares the ratio between the target bandwidth and the previous bandwidth to determine the appropriate row of the Bw Ratio column and then with the appropriate Coef-diff having been determined to select the appropriate column a Delta Index Value is selected. This Delta Bw Index is then added to the previous Bw Index to determine a new BW Index which is then used in the table of FIG. 11 to select the new thresholds. You now exit from the Change bandwidth and Adjust Band Width routine to complete the processing as for the first frame.

Detailed Description Text (94):
However as previously described, despite taking into account the motion characteristic of the image, the new bandwidth value can be far away from the target bandwidth. Thus, frame rate value=Change Frame Rate (target Bw/new Bw). One stores the new Bw in a variable register (previous frame Bw) to be reused by the incoming frame.

Detailed Description Text (95):
Since the frame rate routine is used in the two cases described above in detail, it is also described in detail. The main concern in this routine is when the current bandwidth is larger than the target bandwidth. This means that the bandwidth ratio of target Bw/current Bw is smaller than unity. If the bandwidth ratio<1 the frame rate needs to be modified to reach the target bandwidth.

CLAIMS:

1. A method for compressing input color video data for moving images for a web stream to be transmitted over a transmission channel having a variable bandwidth available to client users having receivers that receive data input at different acceptable data rates, which method comprises the steps of:

compressing an input color video data stream having a determinable frame rate for moving images on the basis of a target bandwidth on a data transmission channel having a variable bandwidth to provide a compressed output video data stream having a determinable frame rate;

transmitting the compressed output color video data stream to at least one receiver over said transmission channel;

determining the actual available bandwidth of the transmission channel on a frame by frame basis as the compressed output video data stream is transmitted to at

least one receiver,

changing the compression of said in put color video data stream in response to comparing the actual available bandwidth on said transmission channel during transmission of the present frame to the available bandwidth on said transmission channel at the time the last frame of said compressed data stream was transmitted thereover whereby the next frame of the subsequent transmission of said output color video data stream is transmitted in packets at a rate commensurate with the actual bandwidth available to a user of said transmission channel at the time of transmission of the previous frame of said data stream,

continually determining the actual available bandwidth of the transmission channel as said compressed output color video data stream is transmitted to at least one receiver, and

continually changing the compression of said output color video data stream in response to a control signal generated from a comparison between the present actual available bandwidth on the data transmission channel and the actual bandwidth on said transmission channel for the immediately previous frame of said data stream whereby the data rate of the packets of said compressed color video data stream is continually adapted to be commensurate with the available variable bandwidth on the transmission channel to dynamically control the packets of said compressed color video data stream, said compression of the input color video data stream includes providing a plurality of levels of subband coding of the input color video data stream, spectrally decomposing selected blocks of pixels of said subband levels to image transform coefficients representative of a level of movement within each pixel block and said control signal is determined by detecting the difference between the present available bandwidth of said output color video data stream and the available bandwidth of the color video data stream transmitted in an immediately previous frame and dynamically changing the rate of compression, as necessary, of said input color video data stream in response to said control signal to provide an output color video data stream in packets having a data rate commensurate with the available bandwidth on the transmission channel to receive such packets.

2. The method of claim 1 which further includes the steps of determining the data rate capability of the receiver of at least one client user to receive said output color video data stream, and

adjusting the data rate of said output color video data stream to be an acceptable data rate for the receiver of the client user.

4. The method of claim 3 wherein the threshold of movement is provided by two thresholds with a first threshold being determined by the difference in intensity level between corresponding image transfer coefficients in two successive frames and if this difference exceeds said first threshold

the image transform coefficient is declared to have moved and with the second threshold being set on the number of image transfer coefficients that have moved within a block of coefficients in a frame, and

the rate of compression of the input color video data stream is controlled by varying the first and second coefficient thresholds in accordance with said control signal to achieve a desired output data rate.

5. The method of claim 3 wherein the data rate of the output color video data stream is further adjusted by adjusting the frame rate of said output data stream on a frame by frame basis whereby said data rate of said output data stream becomes commensurate with the available bandwidth on the transmission channel.

19. A cost effective method for compressing input color video data for moving images for the purpose of streaming a compressed output color video data stream over a packet switched network in the form of a data transmission channel having a variable bandwidth to client users having receivers that provide high quality visual images and which receive data input at different rates, which method comprises the steps of:

compressing an input digital color video data stream having moving images at a predetermined data rate on the basis of a target bandwidth available on a data transmission channel to provide a compressed output color video data stream comprising a plurality of digital video data streams representative of various levels of resolution of a visual image represented in successive video frames and spectrally decomposing the image frames of the input color video signal into a plurality of multilevel frequency subbands and to a plurality of levels and selecting from said frequency subbands a base level digital color video raw data stream that is transmitted in the output video data stream for creation of a base level visual image and at least two additional digital color data streams to provide additional visual details of a visual image as may be desired by a client user;

transmitting the compressed output color video data stream in packets of determinable length to at least one receiver over a data transmission channel;

determining the present bandwidth available on the transmission channel on a frame by frame basis as the compressed output video data streams were transmitted to at least one receiver;

adjusting the rate of compression of said input digital color video data stream in response to such determination on a frame by frame basis whereby the immediately succeeding transmission of a frame of said output digital color video data stream is commensurate with the present bandwidth available on said data transmission channel; and

decompressing the output digital color video data stream at each client user to provide a color visual image having motion.

20. The method of claim 19 wherein the input digital color video data and the output digital color video data stream are presented in successive frames and the adjustment of the rate of compression is made on a frame-by-frame basis in response to a control signal generated by a comparison of the data rate of a color video data stream presently being transmitted for a frame to the immediately previous data rate of the previous frame.

29. The method of claim 19 wherein the rate of compression is also determined by the acceptable data rate of the least capable receiver of a plurality of client users and by a control signal generated by a comparison of the bandwidth of a data stream presently being transmitted for a video frame to the bandwidth on said channel that was available to the immediately preceding video frame of such data stream.

34. Cost effective method for compressing input color video data for moving images for the purpose of streaming a compressed output color video data stream over a packet switched network in the form of a data transmission channel having a variable bandwidth to client users having receivers that provide high quality visual images and which may receive data input at different acceptable rates, which method comprises the steps of:

compressing an input digital color video data stream having moving images at a predetermined data rate on the basis of a target bandwidth available on a data transmission channel to provide a compressed output color video data stream

comprising a plurality of digital video data streams representative of various levels of resolution of a visual image represented in successive video frames and performing scalar quantization on said additional digital color video data streams to generate quantized additional data streams, performing motion detection on said base level color video raw data stream and said quantized additional digital color video data streams and thereafter performing a loss-less compression coding routine on said quantized additional motion data streams to provide said output color video data stream comprising a motion detected base level data stream and motion detected quantized additional digital color video data streams;

transmitting the compressed output color video data stream in packets of determinable length to at least one receiver over a data transmission channel;

determining the present bandwidth available on the transmission channel on a frame by frame basis as the compressed output video data streams were transmitted to at least one receiver;

adjusting the rate of compression of said input digital color video data stream in response to such determination on a frame by frame basis whereby the immediately succeeding transmission of a frame of said output digital color video data stream is commensurate with the present bandwidth available on said data transmission channel; and

decompressing the output digital color video data stream at each client user to provide a color visual image having motion.

Previous Doc Next Doc Go to Doc#

Previous Doc Next Doc Go to Doc#
First Hit Fwd Refs

Generate Collection

L7: Entry 8 of 9 File: USPT Mar 9, 1999

DOCUMENT-IDENTIFIER: US 5881245 A
TITLE: Method and apparatus for transmitting MPEG data at an adaptive data rate

Application Filing Date (1):
19960910

Detailed Description Text (28):
Thus, to reiterate, system interface card 530 is implemented as an interdependent, two processor arrangement (namely, pump microprocessor 708 and control microprocessor 712). The resources of the control microprocessor 712 primarily perform the managing and scheduling functions for delivery of the incoming interleaved MPEG-2 transport stream to the desired ones of the playback cards 504 (i.e., the BOSS functions). The resources of pump microprocessor 708, in conjunction with SCSI-2 fast and wide interface 702, are dedicated to handling the incoming data transfers because of the high throughput rate required and preferably communicate buffer or queue status data, use for rate control, back to the stream server 502. This two processor architecture allows the appropriate data streams to be delivered to a plurality of up to N playback cards 504, thereby effectively spreading the cost of one SCSI drive over a plurality of up to N playback cards 504. This is in contrast to conventional arrangements in which there is a one-to-one correspondence, rather than a one-to-many correspondence, between the intelligence and the decoding functions, that is, normally a single processor is dedicated to each channel or playback card.

Detailed Description Text (39):
For example, consider a simple I/O process in which the stream server 502 is the "initiator" and uses the SCSI bus 532 to transfer a transport stream(s) to a system interface card 530 which is the "target". The I/O process begins with the stream server 502 (or initiator) arbitrating for the SCSI bus 532. Upon winning arbitration, the stream server 502 selects one of the M system interface cards 530 as a target by selecting an appropriate SCSI ID. Although the SCSI protocol supports the selection of one of up to eight (8) targets, due to bandwidth limitations of the SCSI bus data, in a preferred embodiment of the present invention, each SCSI bus has only two (2) associated targets (i.e., system interface cards 530). Once the particular system interface card 530 (or target) is selected, that system interface card 530 (or target) assumes control of the I/O process.

Detailed Description Text (78):
Data regarding the state(s) of the queue(s) of the buffer memory 710 of a system interface card 530 may be provided to the stream server 502 via the pump microprocessor 708, bus 704, and SCSI-2 fast and wide interface 702. Alternatively, such data may be provided to the stream server 502 via control microprocessor 712 over a separate back channel (e.g., Ethernet) link. In either case, the stream server 502 uses such data to adjust the rate at which it provides the transport stream.

Previous Doc Next Doc Go to Doc#

Previous Doc Next Doc Go to Doc#
First Hit Fwd Refs

Generate Collection

L7: Entry 9 of 9 File: USPT Feb 23, 1999

DOCUMENT-IDENTIFIER: US 5875175 A
TITLE: Method and apparatus for time-based download control

Application Filing Date (1):
19970501

Brief Summary Text (26):
According to the invention, other techniques, including grouping packets into temporal sets, may be used by the scheduler to control the flow of transmitted packets to optimize network performance. Temporal sets are groups of packets that the scheduler determines can be transmitted in back-to-back fashion without violating the target bandwidth of any destination. When packets are grouped into temporal sets, the scheduler may set just one alert for a group of packets.

Detailed Description Text (5):
FIGS. 3A and 3B are each diagrams of an ES such as 500b with an adaptor according to alternative embodiments of the invention. FIGS. 3A and 3B both show a host 110 having a CPU process 115, system memory 140, and adaptor 160. Within host 110's processor and address space are higher layer network protocols 120, communicating with adaptor 160 through driver 125. Driver 125 operates in conjunction with a scheduler 130. According to the invention, scheduler 130 determines a schedule for packets to be transmitted based on a target bandwidth for a destination or group of destinations.

Detailed Description Text (6):
Where FIGS. 3A and 3B differ is the manner in which scheduler 130 sets alerts to control packet transmission in order to schedule packets to meet the target transmission bandwidth, as discussed in the section below.

Detailed Description Text (7):
Scheduler 130 determines a time that the next packet for a given destination should be queued on the adaptor for transmission by computing a current time (i.e. the download start time for the present packet to that destination) plus the time it will take to transmit the present packet on the network at the target bandwidth rate known to the scheduler, and then possibly adjusted by values that represent the local transmitter connection transmit time or possibly other local or network intermediate system delay times.

Detailed Description Text (9):
The target destination bandwidth rate may be set according to the actual bandwidth that can be handled by the destination, or may be set for other reasons, such as the control of different types of traffic or the carrying of traffic based on charges levied for access of the network. While the invention is designed for use in networks with variable-width or with fixed-width packets, the invention has particular application when large amounts of data are being transmitted to a destination, and in this case, packets will often be of a maximum (and therefore fixed) size. Therefore, packets of the same size are assumed in some of the following examples.

Detailed Description Text (10):
The operation of a scheduler such as 130 according to the invention will now be explained with reference to the example connections shown in FIGS. 2A-B. For purposes of this discussion, assume that T is a 100 Mbps device on a 100 Mbps link transmitting packets to destinations B and C, each having a target bandwidth of 10 Mbps, and to destination A having a target bandwidth of 30 Mbps. It will be obvious that many other target destinations are possible. If each packet takes 120 microseconds (.mu.s) to transmit over the network at 100 Mbps, and therefore 1200 microseconds at 10 Mbps, according to the invention, if a packet is scheduled for destination B at time 0, the next packet for destination B should be queued by scheduler 130 for download at 0+1200 or 1200 .mu.s. In a different embodiment, the transmit time may be adjusted by a value to take into account various latencies that may be associated with the network or to take advantage of network elasticity. For example, a value such as 120 .mu.s might be subtracted from the queue time for the next packet (perhaps representing the local transmitter time from T to the network) in which case the transmit time for the next packet would be 0 +1200-120, or 1080, microseconds. This would cause the transmitter to slightly "lead" the target bandwidth, but with the assumption that the extra bandwidth can be absorbed by buffering capacity in network intermediate and end system adaptors.

Detailed Description Text (14):
According to a further embodiment of the invention, a scheduler can schedule a number of packets to one, or to a variety of destinations, together in a temporal set. A temporal set is composed of a set of packets which can be transmitted in a back-to-back fashion by the adaptor on its network link without violating the maximum target bandwidth of any destinations.

Detailed Description Text (16):
In one embodiment, for each destination, or group of linked destinations, scheduler 130 determines a desired transmit time for each packet as described above for destination B such that packets to that destination will be transmitted near the target bandwidth of that destination. Scheduler 130 then places as many packets as possible in a temporal set, essentially filling up the set with packets so long as there is a packet available for some destination that could be transmitted immediately after the previous packet (i.e. transmitted back-to-back) without violating the bandwidth of any destination. When there are not more packets ready to be queued that meet that rule, scheduler 130 closes the temporal set and develops an alert for a future time determined by the next packet that needs to be delivered to a destination. Scheduler 130 places packets into their appropriate temporal sets and adjust sets as described below.

CLAIMS:

1. A method for scheduling packets to be downloaded from a host to an adaptor in order to control packet transmission rate comprising:

creating a temporal set of packets that can be transmitted in a back-to-back fashion without violating a maximum bandwidth requirement of said packets' target destination addresses;

characterizing said temporal set by a download start time and an estimated transmission end time;

grouping a packet into a temporal set based on a grouping rule for said packet; and

adjusting said temporal set's download start time and estimated transmission end time according to said grouping rule.

10. A network transmitter capable of controlling its rate of transmission to meet a

target bandwidth comprising:

a destination bandwidth identifier for determining a target bandwidth for a particular destination or group of destinations;

a scheduler for determining when a next packet for a particular destination should be queued for transmit in order to meet said target bandwidth and for programming an alert;

a programmable alert controller for causing transmission of said next packet; and

an interface for receiving packets to be transmitted.

Previous Doc Next Doc Go to Doc#

L7: Entry 4 of 9 File: USPT May 6, 2003

DOCUMENT-IDENTIFIER: US 6560243 B1
TITLE: System and method for receiver based allocation of network bandwidth

Abstract Text (1):
A system receives a flow of data packets via the link and determines a target bandwidth to be allocated to the flow on the link. In response to the flow, the receiving system transmits data to the sending system. The transmitted data control the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth allocated to the flow. In one embodiment, the rate at which the transmitted data from the receiving system arrive at the sending system determines the rate at which the sending system transmits the subsequent data packets. The receiving system can control the rate by delaying its response to the sending system according to a calculated delay factor. In another embodiment, the data transmitted from the receiving system to the sending system indicate a maximum amount of data that the receiving system will accept from the sending system in a subsequent data transmission. The maximum amount is determined so that when the sending system transmits subsequent data packets according to that amount, data is transmitted by the sending system to the receiving system at a rate approximating the target bandwidth.

Application Filing Date (1):
19990430

Brief Summary Text (11):
The present invention relates to methods and systems for allocating bandwidth on a communication link in a network operating according to a flow-controlled protocol where a receiving system receives a flow of data packets from a sending system via the communication link. The receiving system determines a target bandwidth to be allocated to the flow on the link and transmits data to the sending system. The transmitted data provides feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system in response to the transmitted data, data is transmitted by the sending system to the receiving system at a rate approximating the target bandwidth.

Brief Summary Text (13):
In another embodiment of the method, the data transmitted from the receiving system to the sending system indicate a maximum amount of data that the receiving system will accept from the sending system in a subsequent data transmission. The maximum amount is determined so that when the sending system transmits subsequent data packets according to that amount, data is transmitted by the sending system to the receiving system at a rate approximating the target bandwidth.

Brief Summary Text (15):
In still yet another embodiment, the method calculates a period of time for which the receiving system does not transmit data to the sending system. The delay is from when the receiving system would normally transmit the data operating according to the network protocol. The transmission of subsequent data packets on the link by the sending system is thereby delayed for the period of time, causing the data

transmission rate of the sending system to approximate the target bandwidth on the link allocated to the flow. In one embodiment of the invention, the method assigns the flow to a class of flows, and the calculated period of time applies to each flow in that class.

Brief Summary Text (16):
The invention may be implemented in a system including a network interface for receiving a flow of data packets from a sending system on the network via the link. A processor, coupled to the network interface, determines a target bandwidth to be allocated to the flow on the link. The network interface transmits data to the sending system in response to receiving the flow of data packets. The data transmitted by the interface controls the sending system such that the sending system transmits subsequent data packets at a rate corresponding to the target bandwidth allocated to the flow.

Detailed Description Text (7):
Examples of application program 154 are Web browsers that implement rules for data transfer between computers, such as the Hypertext Transfer Protocol (HTTP) or the File Transfer Protocol (FTP). The operating system 158 is a program or set of programs that provides basic services used by the application program 154 that run on the CPU 148. The file system 168 provides program files to operating system 158. The default policy file 170 provides default bandwidth allocation policies for use by the operating system 158 or protocol stack 156 when determining the target bandwidth of data flows. For instance, the default bandwidth allocation policy may automatically classify data flows into various classes based on their duration and actual (monitored) bandwidth usage, and then assign a target bandwidth to each data flow based on its assigned class. The configuration file 172 is used to record specific data flow bandwidth allocations and/or bandwidth allocation policies to be used with specified application programs, and thus to override the default bandwidth allocation policies.

Detailed Description Text (16):
The present invention takes advantage of a fundamental feature of a flow-controlled protocol, which is that during steady state operation the ACK transmission rate of the receiver controls the data transmission rate of the sender. The receiver, then, can use the ACK packets to control the bandwidth allocated to the sender by controlling when to issue the ACK packets and by constraining the amount of data the receiver is willing to receive in a subsequent data transmission. The sender, then, can send more data only after receiving an ACK packet in which the receiver has advanced its advertised window.

Detailed Description Text (22):
The receiver of the invention determines a target bandwidth to be given to the sender. The receiver can make this determination according to a policy that defines how the bandwidth of the bottleneck link is to be shared among various concurrent flows. Given a description of the data flow, such as, for example, the flow's source address, the receiver can determine the fraction of bandwidth, B, allocated to that flow according to governing policy. When the bottleneck link is idle, the full bandwidth of the link should be allocated to the flow.

Detailed Description Text (24):
In one embodiment of the invention, the particular application program 154 that receives the data determines the bandwidth. That is, the program 154 assigns a target bandwidth to each data flow as that flow is created (or at a subsequent time in the existence of that flow) and communicates this target to the operating system 158 or protocol stack 156.

Detailed Description Text (25):
In other embodiments of the invention, the operating system 158 or protocol stack 156 determines the bandwidth directly, based on information provided by a default

policy file 170 in the operating system 154 through an user interface, matched against the particular parameters of the flow (such as the identification of the sender or the protocol type). For example, the default policy file 170 may inform the operating system 158 (or protocol stack 156) that all HTTP flows arriving from a given system X are to be assigned a low target bandwidth BX, and that all FTP flows arriving from a given system Z are to be assigned a high target bandwidth BZ. Changes to existing application program 154 would be unnecessary for such embodiments.

Detailed Description Text (30):
Referring back to FIG. 1, during steady state operation, the bandwidth used by the data flow from the server 106 is approximately the advertised window size of the client 102 divided by the round-trip time. The client 102, then, can achieve the target bandwidth for the flow by adjusting the advertised window size or the round-trip time or both.

Detailed Description Text (32):
In one embodiment of the invention, the client 102 uses the size of the advertised window to control the data transmission rate of the server 106. In the ACK packets 126, the client 102 indicates the amount of data it will accept from the server 102 in a subsequent data packet transmission. This amount is calculated to achieve the target bandwidth when the server 106 transmits subsequent data packets in response to the ACK packets 126.

Detailed Description Text (33):
Through use of advertised windows, the client 102 can either increase or decrease the rate at which the server 106 transmits data in order to achieve the target bandwidth. The client 102 can attain a decrease in rate by shrinking the window size advertised to the server 106. The client 102 can increase the transmission rate of the server 102 by shrinking the window size advertised to other data flow sources that are concurrently using the bottleneck link 118. In this event, the server 106 would gradually discover bandwidth made available by the decreased activity of the other data flow sources.

Detailed Description Text (38):
Delaying the transmission of ACK packets 126 to the server 106 is another technique by which the client 102 can control the server data transmission rate. The delay is imposed from the time when the client 102 would normally transmit the ACK 126, that is, immediately upon receiving the data packet. (For some TCP implementations, the ACK is already delayed so that a single ACK packet acknowledges more than one data packet--the delay of the invention is imposed when that delayed ACK would normally be transmitted. Thus, the acknowledgments of the invention are called "delay-extended" acknowledgments.)

Detailed Description Text (39):
More specifically, the client 102 calculates a delay factor, D, that will cause subsequent data transmissions by the server 106 to be delayed appropriately to achieve the target bandwidth allocated to the flow. Conceptually, the client 102 uses a model function, M(F, T, P), for predicting the bandwidth consumed by a given flow F, at time T, given a set of parameters P. The parameter set P includes the round-trip time between server 106 and client 102, the segment size being used, the size of the server's congestion window, and perhaps other information about the flow F and the state of the network. The parameter set P also includes the calculated delay factor, D, which is the period of time that each acknowledgment for flow F is delayed. That is, if the most recent acknowledgment for flow F was sent by the client 102 M seconds ago, and under normal circumstances the next acknowledgment is due to be transmitted, that next acknowledgment is delayed instead for M*(D-1) seconds.

Detailed Description Text (49):

Successful use of this approach depends on an appropriate division of flows into feedback classes (i.e., classes of flows sharing a feedback-controlled value for D). The following exemplary classes of flows may produce an accurate delay factor for each flow in that class (other classes are possible): One class for each short-lived flow with a different target bandwidth. For example, all short-lived flows with a target bandwidth of B/1000 would be assigned to one class, all short-lived flows with a target bandwidth of B/500 would constitute another class. One class for each long-lived flow. Because a long-lived flow is likely to have a stable steady-state behavior, applying feedback control to derive a value for the delay factor D value should be feasible. One class for each "badly-behaved" flow with a different target bandwidth. A badly-behaved flow is a flow whose observed use of bandwidth significantly exceeds the target bandwidth for its originally assigned class. By isolating badly-behaved flows from well-behaved flows, the bandwidth allocated to badly-behaved flows can be controlled without adversely affecting the well-behaved flows in the same originally-assigned class.

Detailed Description Text (50):
Normally, a data flow would start as a short-lived flow in one of the bandwidth-specific classes. The flow would subsequently terminate or continue long enough to become classified as a long-lived flow. The receiving system would monitor the actual bandwidth consumption of each flow, compare the consumption against the target bandwidth allocated to that flow, and classify that flow as badly-behaved if appropriate.

Detailed Description Text (52):
FIG. 4 shows an exemplary process 200 executed by the client to allocate bandwidth of a bottleneck link to the server 106 according to the invention. In general overview, the client 102 receives a flow of data packets from the server 106 via the bottleneck link (step 204). The client 102 determines a target bandwidth to be allocated to the server 106 for the data flow on the link (step 208).

Detailed Description Text (53):
In some embodiments, the determination of the target bandwidth is based on the application program associated with the data flow. Alternately, the target bandwidth determination may include assigning the data flow to one of a number of classes (222). The assignment of each data flow to a class may be based on a set of default bandwidth allocation policies as well as policies specified in the configuration file 172 (FIG. 2).

Detailed Description Text (55):
At step 214, the client 102 computes a delay factor for delaying the issuing of acknowledgments to the server 106 in response to the flow of data packets. The value of the delay factor is computed such that when the server 106 receives and responds to the acknowledgments and advertised window, the data transmission rate of the server 106 approximates the target bandwidth allocated to the flow. As discussed above, while performing step 214 the process will preferably apply one or more constraints on the rate at which the delay factor may change from one ACK packet to the next.

Detailed Description Text (56):
As discussed above, the target bandwidth for any particular flow may be established either by an application program associated with the flow, or by a set of default policies indicated in a default policy file 170 (FIG. 2).

Detailed Description Text (58):
In another embodiment of the process 200, step 212 is omitted, and data transmission rate of the server 106 is controlled by the rate of the delayed acknowledgments. In yet another embodiment, step 214 is omitted, and the size of the advertised window controls the data transmission rate.

CLAIMS:

1. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving a flow of data packets at a receiving system from a sending system on the network via the link; determining at the receiving system a target bandwidth for the flow on the link; transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; and calculating a period of time for which the receiving system does not transmit data to the sending system.

8. The method of claim 1, wherein the target bandwidth is determined by a network protocol process.

9. The method of claim 1, wherein when the link is idle, the target bandwidth equals a full bandwidth of the link.

10. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving a flow of data packets at a receiving system from a sending system on the network via the link; determining at the receiving system a target bandwidth for the flow on the link; and transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; wherein the flow is initially assigned to a first class of flows, and is subsequently assigned to a second class of flows when the flow continues for longer than a predetermined period of time; the target bandwidth comprising a first target bandwidth for flows in the first class and a second target bandwidth for flows in the second class.

11. The method of claim 10, further comprising the step of: assigning the flow to a third class of flows when the flow uses substantially more bandwidth than the target bandwidth allocated to that flow while in the second class of flows.

12. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving flows of data packets at a receiving system from sending systems on the network via the link; determining at the receiving system a target bandwidth for the flows on the link; transmitting data to the sending systems that will cause the sending systems to transmit subsequent data packets to the receiving system at a rate approximating the target bandwidth for the flows of data packets from the sending systems; and for each flow of at least a subset of the flows, calculating a period of time for which the receiving system does not transmit data to the sending system.

13. A computer system for allocating bandwidth on a communication link in an operating network, comprising: a network interface for receiving a flow of data packets from a sending system on the network via the link; and a processor coupled to the network interface for determining a target bandwidth for the flow on the link, the network interface transmitting data to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that the sending system transmits subsequent data packets at a rate approximating the target bandwidth for the flow when responding to the transmitted data; wherein the processor is configured to calculate a period of time for which the receiving system does not transmit data to the sending system.

18. The computer system of claim 13, wherein the target bandwidth is determined by

an application program that receives the flow of data packets.

19. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving a flow of data packets at a receiving system from a sending system on the network via the link; determining at the receiving system a target bandwidth for the flow on the link; and transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; and calculating and applying a period of time for which the receiving system does not transmit data to a plurality of flows of a class; wherein the target bandwidth is determined by an application program that receives the flow of data packets.

22. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving a flow of data packets at a receiving system from a sending system on the network via the link; determining at the receiving system a target bandwidth for the flow on the link; and transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; wherein the target bandwidth is determined by an application program that receives the flow of data packets; and the flow is initially assigned to a first class of flows, and is subsequently assigned to a second class of flows when the flow continues for longer than a predetermined period of time; the target bandwidth comprising a first target bandwidth for flows in the first class and a second target bandwidth for flows in the second class.

23. The method of claim 22, further comprising the step of: assigning the flow to a third class of flows when the flow uses substantially more bandwidth than the target bandwidth allocated to that flow while in the second class of flows.

24. A method for allocating bandwidth on a communication link in an operating network, comprising the steps of: receiving a flow of data packets at a receiving system from a sending system on the network via the link; determining at the receiving system a target bandwidth for the flow on the link; and transmitting data from the receiving system to the sending system in response to the flow of data packets received, the transmitted data providing feedback to the sending system such that when the sending system transmits subsequent data packets to the receiving system, such subsequent data packets are transmitted at a rate approximating the target bandwidth determined by the receiving system; wherein the target bandwidth is determined by an application program that receives the flow of data packets; and the data transmitted by the receiving system to the sending system includes acknowledgment of receipt of a particular data packet in the flow of data packets.

Previous Doc Next Doc Go to Doc#

Previous Doc Next Doc Go to Doc#
First Hit Fwd Refs

Generate Collection

L23: Entry 1 of 1 File: USPT May 6, 2003

DOCUMENT-IDENTIFIER: US 6560243 B1
TITLE: System and method for receiver based allocation of network bandwidth

Brief Summary Text (15):
In still yet another embodiment, the method calculates a period of time for which the receiving system does not transmit data to the sending system. The delay is from when the receiving system would normally transmit the data operating according to the network protocol. The transmission of subsequent data packets on the link by the sending system is thereby delayed for the period of time, causing the data transmission rate of the sending system to approximate the target bandwidth on the link allocated to the flow. In one embodiment of the invention, the method assigns the flow to a class of flows, and the calculated period of time applies to each flow in that class.

Previous Doc Next Doc Go to Doc#